

Sparsification of Two-Variable Valued CSPs

Arnold Filtser*

Robert Krauthgamer†

September 8, 2015

Abstract

A valued constraint satisfaction problem (VCSP) instance (V, Π, w) is a set of variables V with a set of constraints Π weighted by w . Given a VCSP instance, we are interested in a re-weighted sub-instance $(V, \Pi' \subset \Pi, w')$ such that preserves the value of the given instance (under every assignment to the variables) within factor $1 \pm \epsilon$. A well-studied special case is cut sparsification in graphs, which has found various applications. We show that a VCSP instance consisting of a single boolean predicate $P(x, y)$ (e.g., for cut, $P = \text{XOR}$) can be sparsified into $O(|V|/\epsilon^2)$ constraints if and only if the number of inputs that satisfy P is anything but one (i.e., $|P^{-1}(1)| \neq 1$). Furthermore, this sparsity bound is tight unless P is a relatively trivial predicate. We conclude that also systems of 2SAT (or 2LIN) constraints can be sparsified.

1 Introduction

The seminal work of Benczúr and Karger [BK96] showed that every edge-weighted undirected graph $G = (V, E, w)$ admits cut-sparsification within factor $(1 + \epsilon)$ using $O(\epsilon^{-2}n \log n)$ edges, where we denote throughout $n = |V|$. To state it more precisely, assume that edge-weights are always non-negative and let $\text{Cut}_G(S)$ denote the total weight of edges in G that have exactly one endpoint in S . Then for every such G and $\epsilon \in (0, 1)$, there is a re-weighted subgraph $G_\epsilon = (V, E_\epsilon \subseteq E, w_\epsilon)$ with $|E_\epsilon| \leq O(\epsilon^{-2}n \log n)$ edges, such that

$$\forall S \subset V, \quad \text{Cut}_{G_\epsilon}(S) \in (1 \pm \epsilon) \cdot \text{Cut}_G(S), \quad (1)$$

and moreover, such G_ϵ can be computed efficiently.

This sparsification methodology turned out to be very influential. The original motivation was to speed up algorithms for cut problems – one can compute a cut sparsifier of the input graph and then solve an optimization problem on the sparsifier – and indeed this has been a tremendously effective approach, see e.g. [BK96, BK02, KL02, She09, Mad10]. Another application of this remarkable notion is to reduce space requirement, either when storing the graph or in streaming algorithms [AG09]. In fact, followup work offered several refinements, improvements, and extensions (such as to spectral sparsification or to cuts in hypergraphs, which in turn have more applications) see e.g. [ST04, ST11, SS11, dCHS11, FHHP11, KP12, NR13, BSS14, KK15]. The current bound for

*Ben-Gurion University of the Negev, Israel. Partially supported by the Lynn and William Frankel Center for Computer Sciences. Email: arnoldf@cs.bgu.ac.il

†Weizmann Institute of Science, Israel. Work supported in part by the Israel Science Foundation grant #897/13 and the US-Israel BSF grant #2010418. Email: robert.krauthgamer@weizmann.ac.il

cut sparsification is $O(n/\epsilon^2)$ edges, proved by Batson, Spielman and Srivastava [BSS14], and it is known to be tight [ACK⁺15].

We study the analogous problem of sparsifying Constraint Satisfaction Problems (abbreviated CSPs), which was raised in [KK15, Section 4] and goes as follows. Given a set of constraints on n variables, the goal is to construct a sparse sub-instance, that has approximately the same value as the original instance under *every possible assignment*, see Section 2 for a formal definition. Such sparsification of CSPs can be used to reduce storage space and running time of many algorithms.

We restrict our attention to two-variable constraints (i.e., of arity 2) over boolean domain (i.e. alphabet of size 2). To simplify matters even further we shall start with the case where all the constraints use the same predicate $P : \{0, 1\}^2 \rightarrow \{0, 1\}$. This restricted case of CSP sparsification already generalizes cut-sparsification — simply represent every vertex $v \in V$ by a variable x_v , and every edge $(v, u) \in E$ by the constraint $x_v \neq x_u$.

Observe that such CSPs capture also other interesting graph problems, such as the *uncut edges* (using the predicate $x_v = x_u$), *covered edges* (using the predicate $x_v \vee x_u$) or the *directed-cut edges* (using the predicate $x_v \wedge \neg x_u$). Even though these graph problems are well-known and extensively studied, we are not aware of any sparsification results for them, and at a first glance such sparsification may even seem surprising, because these problems do not have the combinatorial structure exploited by [BK96] (a bound on the number of approximately minimum cuts), or the linear-algebraic description used by [SS11, BSS14] (as quadratic forms over Laplacian matrices).

Results. For CSPs consisting of a single predicate $P : \{0, 1\}^2 \rightarrow \{0, 1\}$, we show in Theorem 3.7 that a $(1 + \epsilon)$ -sparsifier of size $O(n/\epsilon^2)$ always exists if and only if $|P^{-1}(1)| \neq 1$ (i.e., P has 0,2,3 or 4 satisfying inputs). Observe that the latter condition includes the two graphical examples above of uncut edges and covered edges, but excludes directed-cut edges. We further show in Theorem 4.1 that our sparsity bound above is tight, except for some relatively trivial predicates P . We then build on our sparsification result in Section 5 to obtain $(1 + \epsilon)$ -sparsifiers for other CSPs, including 2SAT (which uses 4 predicate types) and 2LIN (which uses 2 predicate types).

Finally, we explore future directions, such as more general predicates and a generalization of the sparsification paradigm to sketching schemes. In particular, we see that the above dichotomy according to number of satisfying inputs to the predicate extends to sketching.

2 Two-Variable Boolean Predicates and Digraphs

A *predicate* is a function $P : \{0, 1\}^2 \rightarrow \{0, 1\}$ (recall we restrict ourselves throughout to two variables and a boolean domain). Given a set of variables V , a *constraint* $\langle (v, u), P \rangle$ consists of a predicate P and an ordered pair (v, u) of variables from V . For an assignment $A : V \rightarrow \{0, 1\}$, we say that A *satisfies* the constraint whenever $P(A(v), A(u)) = 1$. A VCSP (Valued Constraint Satisfaction Problem) instance \mathcal{I} is a triple (V, Π, w) , where V is a set of variables, Π is a set of constraints over V (each of the form $\pi_i = \langle (v_i, u_i), p_i \rangle$), and $w : \Pi \rightarrow \mathbb{R}_+$ is a weight function. The *value* of an assignment $A : V \rightarrow \{0, 1\}$ is the total weight of the satisfied constraints, i.e.,

$$\text{Val}_{\mathcal{I}}(A) := \sum_{\pi_i \in \Pi} w(\pi_i) \cdot p_i(A(v_i), A(u_i)).$$

For $\epsilon \in (0, 1)$, an ϵ -*sparsifier* of \mathcal{I} is a (re-weighted) sub-instance $\mathcal{I}_{\epsilon} = (V, \Pi_{\epsilon} \subseteq \Pi, w_{\epsilon})$ where

$$\forall A : V \rightarrow \{0, 1\}, \quad \text{Val}_{\mathcal{I}_{\epsilon}}(A) \in (1 \pm \epsilon) \cdot \text{Val}_{\mathcal{I}}(A).$$

x_1	x_2	$\vec{0}$	nOr	01	0x	Dicut	x0	Cut	nAnd	And	unCut	$x1$	$\overline{10}$	1x	$\overline{01}$	Or	$\vec{1}$
0	0		1		1		1		1		1		1		1		1
0	1			1	1			1	1			1	1			1	1
1	0					1	1	1	1					1	1	1	1
1	1									1	1	1	1	1	1	1	1

Figure 1: All possible predicates $P : \{0, 1\}^2 \rightarrow \{0, 1\}$, where blank cells denote value 0. Predicates $0x, x0, x1, 1x$ are determined by a single variable. Predicates $01, \text{Dicut}, \overline{10}, \overline{01}$ are satisfied by a single assignment or all but a single one.

The goal is to minimize the number of constraints, i.e., $|\Pi_\epsilon|$. There are 16 different predicates $P : \{0, 1\}^2 \rightarrow \{0, 1\}$, which are listed in Figure 1 with names for easy reference.

We first focus on the case where all the constraints in Π use the same predicate P ,¹ in which case we can represent the VCSP \mathcal{I} by an edge-weighted digraph $G^\mathcal{I} = (V, E, w)$. Each variable in V is represented by a vertex, and each constraint over the pair (v, u) will be represented by a directed edge from v to u , with the same weight as the constraint (formally, $E = \{(v, u) \mid (\langle v, u \rangle, P) \in \Pi\}$, and abusing notation set edge weights $w(v, u) = w(\langle v, u \rangle, P)$). This transformation preserves all the information about the VCSP and allows us to make reductions between VCSPs with different predicates P as their sole predicate.

Given a digraph G , a predicate P and a subset $S \subseteq V$, define

$$P_G(S) := \sum_{(v,u) \in E} P(\mathbf{1}_S(v), \mathbf{1}_S(u)) \cdot w((v, u)),$$

where $\mathbf{1}_S$ denotes the indicator function. For example, applying this definition to the cut predicate $\text{Cut} : (x, y) \rightarrow \mathbf{1}_{\{x \neq y\}}$, we have

$$\text{Cut}_G(S) = \sum_{(v,u) \in E} \text{Cut}(\mathbf{1}_S(v), \mathbf{1}_S(u)) \cdot w((v, u)) = \sum_{(v,u) \in E} |\mathbf{1}_S(v) - \mathbf{1}_S(u)| \cdot w((v, u)),$$

which is just the total weight of the edges crossing the cut S . This matches the definition we gave in the introduction, except for the technical subtlety that G is now a directed graph, which makes no difference for symmetric predicates like Cut . We shall assume henceforth that G is directed.

We shall say that a sub-instance G_ϵ is an ϵ - P -sparsifier of G if

$$\forall S \subseteq V, \quad P_{G_\epsilon}(S) \in (1 \pm \epsilon) \cdot P_G(S).$$

Observe that given an assignment A for the variables V , we can set $S_A := \{u \mid A(u) = 1\}$. It then holds that $\text{Val}_\mathcal{I}(A) = P_{G^\mathcal{I}}(S_A)$, where $G^\mathcal{I}$ is the appropriate digraph for the VCSP. As there a bijection between such VCSPs and digraphs, we conclude

Observation 2.1. *The existence of an ϵ - P -sparsifier $G_\epsilon = (V, E_\epsilon, w_\epsilon)$ for $G^\mathcal{I}$ implies the existence of an ϵ -sparsifier \mathcal{I}_ϵ for \mathcal{I} with $|E_\epsilon|$ constraints.*

¹The collection of predicates used in a VCSP is sometimes called its *signature*. In this paper we mainly deal with VCSPs whose signature is of size one.

Note that the converse is true as well, i.e., an ϵ -sparsifier for \mathcal{I} implies the existence of ϵ -P-sparsifier for $G_{\mathcal{I}}$ of size $|\Pi_{\epsilon}|$. From now on, we focus on finding an ϵ -P-sparsifier for an arbitrary digraph G (for different choices of the predicate P).

3 A Single Predicate

In this section we go over all the predicates $P : \{0, 1\}^2 \rightarrow \{0, 1\}$ and classify them into sparsifiable and non-sparsifiable predicates, see Theorems 3.5, 3.6, and 3.7. For simplicity, we state our sparsification results as existential, but in fact all these sparsifiers can be computed in polynomial time. Our main technique is a simple graph transformation, which seems to be very well-known but in other contexts. We find it surprising that rather different predicates can be analyzed so easily by applying the same elementary transformation.

In our classification, we appeal to two basic predicates, the first of which is **Cut**, which is already known to be sparsifiable.

Theorem 3.1 ([BSS14]). *For every digraph G and parameter $\epsilon \in (0, 1)$, there is an ϵ -Cut-sparsifier for G with $O(|V|/\epsilon^2)$ edges.*

Our second basic predicate is the predicate **And**, which behaves significantly different. We call a digraph $G = (V, E)$ *strongly asymmetric* if for every $(v, u) \in E$ it holds that $(u, v) \notin E$.

Theorem 3.2. *For every strongly asymmetric digraph $G = (V, E, w)$ with strictly positive weights and $\epsilon \in (0, 1)$, every ϵ -And-sparsifier $G_{\epsilon} = (V, E_{\epsilon}, w_{\epsilon})$ must satisfy $E_{\epsilon} = E$.*

Proof. Let $G_{\epsilon} = (V, E_{\epsilon}, w_{\epsilon})$ be such a sparsifier, i.e., for every $S \subseteq V$ it holds that $\text{And}_{G_{\epsilon}}(S) \in (1 \pm \epsilon) \cdot \text{And}_G(S)$. Then for every $e = (v, u) \in E$ we must have $(v, u) \in E_{\epsilon}$, as otherwise for the set $S = \{u, v\}$ it will hold that $\text{And}_{G_{\epsilon}}(\{u, v\}) = 0$ while $\text{And}_G(\{u, v\}) = w(e) > 0$, a contradiction. \square

Remark 3.3. *For every digraph (which is not necessarily strongly asymmetric), the same proof shows that $|E_{\epsilon}| \geq \frac{1}{2}|E|$.*

Remark 3.4. *Our definition of an ϵ -P-sparsifier requires G_{ϵ} to be a subgraph of G , but we can state Theorem 3.2 in a more general way: For every digraph $G_{\epsilon} = (V, E_{\epsilon}, w_{\epsilon})$ (not necessarily a subgraph) such that every $S \subseteq V$ satisfies $\text{And}_{G_{\epsilon}}(S) \in (1 \pm \epsilon) \cdot \text{And}_G(S)$ necessarily E_{ϵ} agrees with E up to the directions of the edges.*

Next, we show that every other predicates is similar either to **Cut** or to **And** in terms of sparsifiability. We describe a reduction that will be useful to show both sparsifiability and non-sparsifiability. (This reduction is based on a well-known transformation of a given graph, called the “bipartite double cover”, see e.g. [BHM80], although we are not aware of its use in the same way.) Let γ be a function that maps a digraph $G = (V, E, w)$ where $V = \{v_1, v_2, \dots, v_n\}$ to a digraph $\gamma(G) = (V^{\gamma}, E^{\gamma}, w^{\gamma})$ where $V^{\gamma} = \{v_{-n}, \dots, v_{-1}, v_1, \dots, v_n\}$, $E^{\gamma} = \{(v_i, v_{-j}) \mid (v_i, v_j) \in E\}$, $w^{\gamma}((v_i, v_{-j})) = w((v_i, v_j))$. For every subset $S \subseteq V$, we introduce the notation $-S := \{v_{-i} \mid v_i \in S\}$, $\bar{S} := \{v_i \mid v_i \in V \setminus S\}$ and $-\bar{S} := \{v_{-i} \mid v_i \in V \setminus S\}$. Figure 2 illustrates the effect of γ on an arbitrary set S .

Theorem 3.5. *For every digraph $G = (V, E, w)$ and $\epsilon \in (0, 1)$ there is a sub-digraph G_{ϵ} with $O(|V|/\epsilon^2)$ edges, such that for every predicate $P \in \{\text{Cut}, \text{unCut}, \text{Or}, \text{nAnd}, \overline{10}, \overline{01}, x0, x1, 0x, 1x, \overline{1}, \overline{0}\}$, the digraph G_{ϵ} is an ϵ -P-sparsifier of G . (Note that G_{ϵ} does not depend on P .)*




Figure 2: The mapping γ applied on G and its effect on an arbitrary $S \subseteq V$. For example, an edge from $v_i \in S$ to $v_j \in \bar{S}$ is represented by an arrow of type 3, and becomes in $\gamma(G)$ an edge from $v_i \in S$ to $v_{-j} \in -\bar{S}$.

Proof. Given G and ϵ , first construct $\gamma(G)$ as above. Next, apply [Theorem 3.1](#) to obtain for $\gamma(G)$ a cut-sparsifier $\gamma(G)_\epsilon = (V^\gamma, E_\epsilon^\gamma \subseteq E^\epsilon, w_\epsilon^\gamma)$, which contains $O(|V^\gamma|/\epsilon^2) = O(|V|/\epsilon^2)$ edges. Now construct a digraph $G_\epsilon = (V, E_\epsilon, w_\epsilon)$ where $E_\epsilon = \{(v_i, v_j) \mid (v_i, v_{-j}) \in E_\epsilon^\gamma\}$ and $w_\epsilon(v_i, v_j) = w_\epsilon^\gamma(v_i, v_{-j})$. Observe that $\gamma(G_\epsilon) = \gamma(G)_\epsilon$, i.e. if we apply γ on G_ϵ we get exactly $\gamma(G)_\epsilon$.

Now suppose that for a predicate P , there is a function $f_P : 2^V \rightarrow 2^{V^\gamma}$ such that for every digraph H on the vertex set V , it holds that

$$\forall S \subset V, \quad P_H(S) = \text{Cut}_{\gamma(H)}(f_P(S)). \quad (2)$$

Then we could apply (2) twice, first to G_ϵ and then to G , and obtain that

$$\forall S \subset V, \quad P_{G_\epsilon}(S) = \text{Cut}_{\gamma(G)_\epsilon}(f_P(S)) \in (1 \pm \epsilon) \cdot \text{Cut}_{\gamma(G)}(f_P(S)) = (1 \pm \epsilon) \cdot P_G(S).$$

Hence, the existence of such a function f_P implies that G_ϵ is an ϵ - P -sparsifier. And indeed, we can show such f_P for some predicates P , as follows.

- $f_{\text{unCut}}(S) = S \cup \bar{S}$;
- $f_{\text{Cut}}(S) = S \cup -S$;
- $f_{0x}(S) = \bar{S}$;
- $f_{x0}(S) = -\bar{S}$;
- $f_{x1}(S) = -S$;
- $f_{1x}(S) = S$;
- $f_{\bar{1}}(S) = S \cup \bar{S}$; and
- $f_{\bar{0}}(S) = \emptyset$.

To verify that $f_{\text{unCut}}(S) = S \cup \bar{S}$ satisfies [Equation 2](#), i.e., that $\text{unCut}_H(S) = \text{Cut}_{\gamma(H)}(S \cup \bar{S})$, observe that both sides consist exactly of the edges of types 1 and 2 in [Figure 2](#). The other predicates can be easily verified similarly, which completes the proof for all $P \in \{\text{Cut}, \text{unCut}, 0x, x0, x1, 1x, \bar{1}, \bar{0}\}$.

To show that G_ϵ is a sparsifier also for predicates $P \in \{\text{Or}, \text{nAnd}, \bar{1}\bar{0}, \bar{0}\bar{1}\}$ we need a slightly more general argument. Suppose that for a predicate P , there are functions $f_P^1, f_P^2, f_P^3 : 2^V \rightarrow 2^{V^\gamma}$ such that for every digraph H on the vertex set V ,

$$P_H(S) = \frac{1}{2} [\text{Cut}_{\gamma(H)}(f_P^1(S)) + \text{Cut}_{\gamma(H)}(f_P^2(S)) + \text{Cut}_{\gamma(H)}(f_P^3(S))]. \quad (3)$$

Then we could apply (3) twice, first to G_ϵ and then to G , and obtain that

$$\begin{aligned} P_{G_\epsilon}(S) &= \frac{1}{2} [\text{Cut}_{\gamma(G)_\epsilon}(f_P^1(S)) + \text{Cut}_{\gamma(G)_\epsilon}(f_P^2(S)) + \text{Cut}_{\gamma(G)_\epsilon}(f_P^3(S))] \\ &\in (1 \pm \epsilon) \cdot \frac{1}{2} [\text{Cut}_{\gamma(G)}(f_P^1(S)) + \text{Cut}_{\gamma(G)}(f_P^2(S)) + \text{Cut}_{\gamma(G)}(f_P^3(S))] \\ &= (1 \pm \epsilon) \cdot P_G(S). \end{aligned}$$

Hence, the existence of such three functions will imply that G_ϵ is an ϵ -P-sparsifier. And indeed, we let

- $f_{\text{Or}}^1(S) = S, f_{\text{Or}}^2(S) = -S, f_{\text{Or}}^3(S) = S \cup -S;$
- $f_{\text{nAnd}}^1(S) = \bar{S}, f_{\text{nAnd}}^2(S) = -\bar{S}, f_{\text{nAnd}}^3(S) = \bar{S} \cup -\bar{S};$
- $f_{\text{IO}}^1(S) = \bar{S}, f_{\text{IO}}^2(S) = -S, f_{\text{IO}}^3(S) = \bar{S} \cup -S;$ and
- $f_{\text{OI}}^1(S) = S, f_{\text{OI}}^2(S) = -\bar{S}, f_{\text{OI}}^3(S) = S \cup -\bar{S}.$

To verify that $f_{\text{Or}}^1, f_{\text{Or}}^2, f_{\text{Or}}^3$ satisfies Equation 3, observe that both sides consist exactly of the edges of types 1, 3, 4 in Figure 2. The other predicates can be easily verified similarly, which completes the proof for all $P \in \{\text{Or}, \text{nAnd}, \text{IO}, \text{OI}\}$. \square

Next, we use γ for a reductions from And to all the remaining predicates. In particular it will imply their “resistance to sparsification”.

Theorem 3.6. *Given parameters n and $m \leq \binom{n}{2}$, there is a digraph $G = (V, E, w)$ with $2n$ vertices and m edges such that for every $\epsilon \in (0, 1)$ and every predicate $P \in \{\text{nOr}, \text{OI}, \text{Dicut}, \text{And}\}$, for every ϵ -P-sparsifier $G_\epsilon = (V, E_\epsilon, w_\epsilon)$ of G it holds that $E_\epsilon = E$. (Note that G does not depend on P .)*

Proof. Let $G = (V, E, w)$ be an arbitrary strongly asymmetric digraph with n vertices, m edges and strictly positive weights. Let $\gamma(G)$ be the digraph constructed by our reduction. Note that $\gamma(G)$ consist of $2n$ vertices and m edges. $\gamma(G)$ will be the digraph for which we will prove the theorem.

Fix some predicate P . Let $\gamma(G)_\epsilon = (V^\gamma, E_\epsilon^\gamma \subseteq E^\epsilon, w_\epsilon^\gamma)$ be some ϵ -P-sparsifier for $\gamma(G)$. Let $G_\epsilon = (V, E_\epsilon, w_\epsilon)$ be a digraph where $E_\epsilon = \{(v_i, v_j) \mid (v_i, v_{-j}) \in E_\epsilon^\gamma\}$ and $w_\epsilon((v_i, v_j)) = w_\epsilon^\gamma((v_i, v_{-j}))$. Note that $\gamma(G_\epsilon) = \gamma(G)_\epsilon$.

Now suppose that there is a function $f_P : 2^V \rightarrow 2^{V^\gamma}$ such that for every digraph H on the vertex set V , it holds that

$$\forall S \subset V, \quad \text{And}_H(S) = P_{\gamma(H)}(f_P(S)). \quad (4)$$

Then we could apply (4) twice, first to G_ϵ and then to G , and obtain that

$$\forall S \subset V, \quad \text{And}_{G_\epsilon}(S) = P_{\gamma(G)_\epsilon}(f_P(S)) \in (1 \pm \epsilon) \cdot P_{\gamma(G)}(f_P(S)) = (1 \pm \epsilon) \cdot \text{And}_G(S).$$

Hence, assuming such a function f exists, G_ϵ is an ϵ -And-sparsifier for G . According to Theorem 3.2, necessarily $E_\epsilon = E$, and in particular $E_\epsilon^\gamma = E^\gamma$.

Hence, The existence of such functions f_P for all $P \in \{\text{nOr}, \text{OI}, \text{Dicut}, \text{And}\}$ will imply our theorem. And indeed, we let

- $f_{\text{And}}(S) = S \cup -S;$
- $f_{\text{nOr}}(S) = \bar{S} \cup -\bar{S};$

- $f_{Dicut}(S) = S \cup -\bar{S}$; and
- $f_{01}(S) = \bar{S} \cup -S$.

To verify that $f_{Dicut}(S) = S \cup -\bar{S}$ satisfies Equation 4, observe that both sides consist exactly of the edges of type 1 in Figure 2. The other predicates can be easily verified similarly. \square

We conclude our main theorem, which basically puts together Theorems 3.5 and 3.6.

Theorem 3.7. *Let P be a binary predicate, and let $\epsilon \in (0, 1)$ be some parameter.*

- *If P has a single “1” in its truth table then there exist a VCSP $\mathcal{I} = (V, \Pi, w)$ with a single predicate P , such that every ϵ - P -sparsifier of \mathcal{I} will have $\Omega(|V|^2)$ constraints.*
- *If P does not has a single “1” in its truth table then for every VCSP $\mathcal{I} = (V, \Pi, w)$ with single predicate P , there exists an ϵ - P -sparsifier with $O(|V|/\epsilon^2)$ constraints.*

4 Lower Bounds (for a Single Predicate)

In this section we will show that Theorem 3.5 is tight. More precisely, we will show that for every $P \in \{\text{Cut}, \text{unCut}, \text{Or}, \text{nAnd}, \bar{1}\bar{0}, \bar{0}\bar{1}\}$, there exists an n -vertex graph G such that every ϵ - P -sparsifier G_ϵ of G must contain $\Omega(n/\epsilon^2)$ edges.² The first step was done by [ACK⁺15], who showed that Theorem 3.1 is tight, i.e., for every n and $\epsilon \in (1/\sqrt{n}, 1)$, there exists n -vertex graph G such that every ϵ -Cut-sparsifier G_ϵ of G must contain $\Omega(n/\epsilon^2)$ edges. Using our reduction γ in similar manner to Theorem 3.5, this lower bound can be extended to unCut based on the fact that $\text{Cut}_G(S) = \text{unCut}_{\gamma(G)}(S \cup -\bar{S})$. However, γ fails to extend the lower bound to predicates with three 1’s in their truth table. To this end, we will define sketching schemes, a variation of sparsification where the goal is to maintain the approximate value of every assignment using a small data structure, possibly without any combinatorial structure, see definition below. We will use a lower bound on the sketch size of Cut from [ACK⁺15] to prove lower bound on the number of edges in a sparsifier (and also on the sketch size) for OR. The extension to other predicates with three 1’s in their truth table is straightforward using γ . Sketching is interesting for its own, and we have further discussion and lower bounds regarding sketching in Section 6.3.

Formally, a *sketching scheme* (or a *sketch* in short) is a pair of algorithms (sk, est). Given a weighted digraph $G = (V, E, w)$ and a predicate P , algorithm sk returns a string sk_G (intuitively, a short encoding of the instance). Given sk_G and a subset $S \subseteq V$, algorithm est returns a value (without looking at G) that estimates $P_G(S)$. We say that it is an ϵ - P -*sketching-scheme* if for every digraph G , and for every subset $S \subseteq V$, $\text{est}(\text{sk}_G, S) \in (1 \pm \epsilon) \cdot P_G(S)$. The *sketch-size* is $\max_G |\text{sk}_G|$, the maximal length of the encoding string over all the digraphs with n variables, often measured in bits. sk might be probabilistic algorithm, but for our purposes it is enough to think only on the deterministic case. Note that an algorithm for constructing ϵ -sparsifiers always provides an ϵ -sketching-scheme, where the sketch-size is asymptotically equal to the number of constraints in the constructed sparsifiers when measured in machine words (and up to logarithmic factors when measured in bits). Sparsification is advantageous over general sketching as it preserves the combinatorial structure of the problem. Nevertheless, one may be interested in constructing sketches as they may potentially require significantly smaller storage.

²The other predicates $\{x0, x1, 0x, 1x, \bar{1}, \bar{0}\}$, are kind of trivial in the sense of sparsification. $\bar{0}$ sparsified by the empty graph. $\bar{1}$ can be sparsified using a single edge. $\{x0, x1, 0x, 1x\}$ could be sparsified using n edges.

Theorem 4.1. Fix a predicate $P \in \{\text{Cut}, \text{unCut}, \text{Or}, \text{nAnd}, \overline{10}\}$, an integer n and $\epsilon \in (1/\sqrt{n}, 1)$. The sketch-size of every ϵ - P -sketching-scheme on n variables is $\Omega(n/\epsilon^2)$. Moreover, there is an n -vertex digraph G , such that every ϵ - P -sparsifier of G has $\Omega(n/\epsilon^2)$ edges.

Proof. We follow the line-of-proof of Theorems 4.1 and 4.2 in [ACK⁺15]. Specifically, they show that the sketch-size of every ϵ -Cut-sketching-scheme is $\Omega(n/\epsilon^2)$ bits, by proving that a certain family \mathcal{F} of n -vertex graphs is hard to sketch, and consequently to sparsify. By similar arguments to Theorem 3.5, this lower bound easily extends to unCut. Indeed, recall that $\text{Cut}_G(S) = \text{unCut}_{\gamma(G)}(S \cup -\bar{S})$, and thus a ϵ -unCut-sparsifier (or sketch) for $\gamma(G)$ yields an ϵ -Cut-sparsifier (or sketch) for G with the same number of edges (size).

Once we prove the lower bound for predicate OR, a reduction from OR using γ will extend it also to nAnd, $\overline{10}$ and $\overline{01}$, because

$$\text{Or}_G(S) = \text{nAnd}_{\gamma(G)}(\bar{S} \cup -\bar{S}) = \overline{01}_{\gamma(G)}(S \cup -\bar{S}) = \overline{10}_{\gamma(G)}(\bar{S} \cup -S). \quad (5)$$

We will thus focus on the predicate OR. As it is symmetric predicate, we can work with graphs rather than digraphs. The main observation in our proof is that for every undirected graph $G = (V, E, w)$, if $\deg_G(v)$ denotes the degree of vertex v , then

$$\forall S \subset V, \quad \text{Cut}_G(S) = 2 \cdot \text{OR}_G(S) - \sum_{v \in S} \deg_G(v). \quad (6)$$

The graph family \mathcal{F} consists of graphs G constructed as follows. Let $s_1, \dots, s_{n/2} \in \{0, 1\}^{1/\epsilon^2}$ be balanced $1/\epsilon^2$ bit-strings (i.e., each s_i has normalized Hamming weight exactly $1/2$), and let the graph G be a disjoint union of the graphs $\{G_j \mid j \in [n/2]\}$, where each G_j is a bipartite graph, whose two sides, each of size $1/\epsilon^2$, are denoted $L(G_j)$ and $R(G_j)$. The edges of G are determined by $s_1, \dots, s_{n/2}$, where each bit string s_i indicates the adjacency between vertex $i \in \cup_j L(G_j)$ and the vertices in the respective $R(G_j)$. They further observe (in Theorem 4.2) that the lower bound holds even if the sketching scheme is relaxed as follows:

1. The estimation is required only for cut queries contained in a single G_j , namely, cut queries $S \cup T$ where $S \subset L(G_j)$ and $T \subset R(G_j)$ for the same j .
2. The estimation achieves additive error μ/ϵ^3 , where $\mu = 10^{-4}$ (instead of multiplicative error $1 \pm \epsilon$).

To prove a sketch-size lower bound for a $(\mu\epsilon)$ -OR-sketching-scheme $(\text{sk}^{\text{OR}}, \text{est}^{\text{OR}})$, we assume it has sketch-size $s = s(n, \epsilon)$ bits, and use it to construct a Cut-sketching-scheme $(\text{sk}^{\text{Cut}}, \text{est}^{\text{Cut}})$ that achieves the estimation properties 1 and 2 on graphs of the aforementioned form, and has sketch-size $s + 2n \log(1/\epsilon)$ bits. Then by [ACK⁺15], this sketch-size must be $\Omega(n/\epsilon^2)$, and we conclude that $s = \Omega(n/\epsilon^2)$ as required.

Given a graph $G \in \mathcal{F}$, let sk_G^{Cut} be a concatenation of sk_G^{OR} and a list of all vertex degrees in G . The degrees in G are bounded by $1/\epsilon^2$, hence the size of sk_G^{Cut} is indeed $s + 2n \log(1/\epsilon)$ bits. Given a cut query $S \cup T$ contained in some G_j , define the estimation algorithm (which we now construct for Cut) to be

$$\text{est}^{\text{Cut}}(\text{sk}_G^{\text{Cut}}, S \cup T) := 2 \cdot \text{est}^{\text{OR}}(\text{sk}_G^{\text{OR}}, S \cup T) - \sum_{v \in S \cup T} \deg_G(v). \quad (7)$$

Let us analyze the error of this estimate. First, observe that as in each G_j there are precisely $\frac{1}{2\epsilon^4}$ edges, $\text{OR}_G(S \cup T) \leq \frac{1}{2\epsilon^4}$, and thus

$$\text{est}^{\text{OR}}(\text{sk}_G^{\text{OR}}, S \cup T) \in (1 \pm \mu\epsilon) \cdot \text{OR}_G(S \cup T) \subseteq \text{OR}_G(S \cup T) \pm \frac{\mu}{2\epsilon^3}.$$

Plugging this estimate into (7) and then recalling our initial observation (6), we obtain as desired

$$\begin{aligned} \text{est}^{\text{Cut}}(\text{sk}_G^{\text{Cut}}, S \cup T) &\in 2 \cdot \text{OR}_G(S \cup T) \pm \frac{\mu}{\epsilon^3} - \sum_{v \in S \cup T} \deg_G(v) \\ &= \text{Cut}_G(S \cup T) \pm \frac{\mu}{\epsilon^3}. \end{aligned}$$

To prove a lower bound on the size of an OR-sparsifier, we follow the argument in [ACK⁺15, Theorem 4.2], which shows that given an ϵ -Cut-sparsifier G_ϵ with $s = s(n, \epsilon)$ edges for a graph $G \in \mathcal{F}$, there is a Cut-sparsifier G_μ of G_ϵ , with additive error $\mu/2\epsilon^3$, such that G_μ has only integer weights and henceforth can be encoded using $O(s(\mu^{-2} + \log(\epsilon^{-2}n/s)))$ bits. In fact, there is nothing special here about Cut. The same proof will work (with the same properties) for predicate OR, assuming a sparsifier is required to be a subgraph (to remove this restriction, just erase all the edges between G_j to G_i for $i \neq j$, which adds only a small additive error).

Now suppose that every graph G of the form specified above admits a $\frac{\mu}{2}\epsilon$ -OR-sparsifier G_ϵ with s edges. Then as explained above (about repeating the argument of [ACK⁺15]) there is a graph G_μ that sparsifies G_ϵ with additive error $\mu/2\epsilon^3$, and can be encoded by a string \mathcal{I}_G of size $O(s \log(\epsilon^{-2}n/s))$ bits (recall that μ is a constant). Use it to construct a Cut-sketching-scheme with additive error μ/ϵ^3 as follows. Given the graph G , set sk_G^{Cut} to be the concatenation of \mathcal{I}_G and a list of the degrees of all the vertices in G . Then $|\mathcal{I}_G| = O(s \log(\epsilon^{-2}n/s)) + 2n \log(1/\epsilon)$. For a cut query $S \cup T$ contained in some G_j , define the estimation algorithm (using the OR sparsifier) to be

$$\text{est}^{\text{Cut}}(\text{sk}_G^{\text{Cut}}, S \cup T) := 2 \cdot \text{OR}_{G_\mu}(S \cup T) - \sum_{v \in S \cup T} \deg_G(v).$$

Then we can again analyze it by plugging the above error bounds and then using (6),

$$\begin{aligned} \text{est}^{\text{Cut}}(\text{sk}_G^{\text{Cut}}, S \cup T) &\in 2 \cdot \text{OR}_{G_\epsilon}(S \cup T) \pm \frac{\mu}{2\epsilon^3} - \sum_{v \in S \cup T} \deg_G(v) \\ &\in 2 \cdot \text{OR}_G(S \cup T) \pm \frac{\mu}{\epsilon^3} - \sum_{v \in S \cup T} \deg_G(v) \\ &= \text{Cut}_G(S \cup T) \pm \frac{\mu}{\epsilon^3}. \end{aligned}$$

By [ACK⁺15], the sketch-size must be $|\mathcal{I}_G| = \Omega(n/\epsilon^2)$, hence $s = \Omega(n/\epsilon^2)$ (for at least one graph $G \in \mathcal{F}$) as required. \square

5 Multiple Predicates and Applications

In this section we extend Theorem 3.5 to VCSPs using multiple types of predicates. In particular, we prove sparsifiability for some classical problems. Again, our sparsification results are stated as existential bounds, but these sparsifiers can actually be computed in polynomial time.

Theorem 5.1. *For every $\epsilon \in (0, 1)$ and a VCSP (V, Π, w) whose constraints $\langle (v, u), P \rangle \in \Pi$ all satisfy $P \notin \{\text{nOr}, \text{01}, \text{Dicut}, \text{And}\}$, there exists an ϵ -sparsifier for \mathcal{I} with $O(|V|/\epsilon^2)$ constraints.*

This bound is tight, according to [Theorem 4.1](#). We prove it by a straightforward application of [Theorem 3.5](#). Partition \mathcal{I} to disjoint VCSPs according to the predicates in the constraints, and then for each sub-VCSP find an ϵ -sparsifier using [Theorem 3.5](#). The union of this sparsifiers is an ϵ -sparsifier for \mathcal{I} . A formal proof follows.

Proof of Theorem 5.1. For each predicate P , let $\Pi^P = \{\pi \in \Pi \mid \pi = \langle (v, u), P \rangle\}$. Note that $\{\Pi^P\}$ forms a partition of Π . For each P , let $\mathcal{I}^P = (V, \Pi^P, w^P)$ where w^P is the restriction of w to Π^P . Let $\mathcal{I}_\epsilon^P = (V, \Pi_\epsilon^P, w_\epsilon^P)$ be an ϵ - P -sparsifier for \mathcal{I}^P with $|\Pi_\epsilon^P| = O(|V|/\epsilon^2)$ constraints according to [Theorem 3.5](#) (recall that $P \notin \{\text{nOr}, \text{01}, \text{Dicut}, \text{And}\}$). Set $\mathcal{I}_\epsilon = (V, \Pi_\epsilon, w_\epsilon)$, $\Pi_\epsilon = \bigcup_P \Pi_\epsilon^P$ and $w_\epsilon = \bigcup_P w_\epsilon^P$. For every assignment A ,

$$\begin{aligned} \text{Val}_{\mathcal{I}_\epsilon}(A) &= \sum_{\pi_i \in \Pi_\epsilon} w_\epsilon(\pi_i) \cdot p_i(A(v_i), A(u_i)) \\ &= \sum_P \sum_{\pi_i \in \Pi_\epsilon^P} w_\epsilon^P(\pi_i) \cdot P(A(v_i), A(u_i)) \\ &\in (1 \pm \epsilon) \cdot \sum_P \sum_{\pi_i \in \Pi^P} w^P(\pi_i) \cdot P(A(v_i), A(u_i)) \\ &= (1 \pm \epsilon) \cdot \sum_{\pi_i \in \Pi} w(\pi_i) \cdot p_i(A(v_i), A(u_i)) \\ &= (1 \pm \epsilon) \cdot \text{Val}_{\mathcal{I}}(A), \end{aligned}$$

and note that indeed $|\Pi_\epsilon| \leq O(n/\epsilon^2)$. □

2SAT (boolean satisfiability problem over constraints with 2 variables) can be viewed as a VCSP which uses only the predicates Or , nAnd , $\overline{10}$ and $\overline{01}$. By [Theorem 5.1](#), for every 2SAT formula Φ over n variables, and for every $\epsilon \in (0, 1)$, there is a sub-formula Φ_ϵ with $O(n/\epsilon^2)$ clauses, such that Φ and Φ_ϵ have the same value for every assignment up to factor $1 + \epsilon$.³

2LIN is a system of linear equations (modulo 2), where each equation contains 2 variables and has a nonnegative weight. Notice that the equation $x + y = 1$ is a constraint using the Cut predicate while the equation $x + y = 0$ is a constraint using the unCut predicate. By [Theorem 5.1](#), if n denotes the number of variables, then for every $\epsilon \in (0, 1)$ we can construct a sparsifier with only $O(n/\epsilon^2)$ equations (i.e., a re-weighted subset of equations, such that on every assignment it agrees with the original system up to factor $1 + \epsilon$).

We note that by our lower bound ([Theorem 4.1](#)), there are instances of 2SAT (2LIN) for which every ϵ -sparsifier must contain $\Omega(n/\epsilon^2)$ clauses (equations).

6 Further Directions

Based on the past experience of cut sparsification in graphs – which has been extremely successful in terms of techniques, applications, extensions and mathematical connections – we expect VCSP

³We use here the version of 2SAT where each clause has weight and every assignment has value rather than the version when we only ask whether there is an assignment that satisfies all the clauses.

sparsification to have many benefits. A challenging direction is to identify which predicates admit sparsification, and our results make the first strides in this direction.

We now discuss potential extensions to our results in the previous sections (which characterize two-variable predicates over a boolean alphabet). We first consider predicates with more variables, and in particular show sparsification for k -SAT formulas, in Section 6.1. We then consider predicates with large alphabets in Section 6.2, showing in particular a sparsifier construction for k -Cut, and that linear equations (modulo $k \geq 3$) are not sparsifiable. We also consider sketching schemes, notable we discuss a more loose sketching model called *for-each* in Section 6.3. Finally, we study *spectral* sparsification for unCut, a notion that preserves some algebraic properties in addition to the “uncuts” in Section 6.4.

6.1 Predicates over more variables and k -SAT

It is natural to ask for the best bounds on the size of ϵ -P-sparsifiers for different predicates $P : \{0, 1\}^k \rightarrow \{0, 1\}$. A first step towards answering this question was already done by [KK15].

Theorem 6.1 ([KK15]). *For every hypergraph $H = (V, E, w)$ with hyperedges containing at most r vertices, and $\epsilon \in (0, 1)$, there is a re-weighted subhypergraph H_ϵ with $O(n(r + \log n)/\epsilon^2)$ hyperedges such that*

$$\forall S \subseteq V, \quad \text{Cut}_{H_\epsilon}(S) \in (1 \pm \epsilon) \cdot \text{Cut}_H(S).$$

Here we say that a hyperedge e is *cut* by S if $S \cap e \neq \emptyset, e$ (i.e., not all the vertices in e are in the same side). Observe that Cut is equivalent to the predicate NAE (not all equal). In particular Theorem 6.1 implies that for every VCSP using only NAE, there is an ϵ -sparsifier with $O(n(r + \log n)/\epsilon^2)$ constraints.

A k -SAT is essentially a VCSP that uses only predicates with a single 0 in their truth table. [KK15] use Theorem 6.1 to construct an ϵ -sketching-scheme with sketch-size $\tilde{O}(nk/\epsilon^2)$ for k -SAT formulas (i.e., only for VCSPs of this particular form). We observe that their sketching scheme can be further used to construct an ϵ -sparsifiers, as follows.

First, recall how the sketching scheme of [KK15] works. Given a k -SAT formula $\Phi = (V, \mathcal{C}, w)$ (variables, clauses, weight over \mathcal{C}), construct a hypergraph H on vertex set $V \cup -V \cup \{f\}$. We associate the literal v_i with vertex v_i , the literal $\neg v_i$ with vertex v_{-i} , and use f to represent the “false”. Each clause becomes a hyperedge consisting of f and (the vertices associated with) the literals in \mathcal{C} (for example $v_5 \vee \neg v_7 \vee v_{12}$ becomes $\{f, v_5, v_{-7}, v_{12}\}$). Observe that given a truth assignment $A : V \rightarrow \{0, 1\}$, if we define $S_A := \{u \mid A(u) = 0\}$, then $\text{Val}_\Phi(A) = \text{Cut}_H(S_A \cup \{f\})$, and using Theorem 6.1 this provides a sketching scheme. Moreover, given an ϵ -Cut-sparsifier H_ϵ for H , let Φ_ϵ be the formula which has only the clauses associated with edges that “survived” the sparsification, with the same weight. Notice that for every assignment A ,

$$\text{Val}_{\Phi_\epsilon}(A) = \text{Cut}_{H_\epsilon}(S_A \cup \{f\}) \in (1 \pm \epsilon) \cdot \text{Cut}_H(S_A \cup \{f\}) = (1 \pm \epsilon) \cdot \text{Val}_\Phi(A) .$$

Theorem 6.2. *Given k -SAT formula Φ over n variables and parameter $\epsilon \in (0, 1)$, there is an ϵ -sparsifier sub-formula ϕ_ϵ with $O(n(k + \log n)/\epsilon^2)$ clauses.*

In contrast, we are not aware of any nontrivial sparsification result for the parity predicate (on $k \geq 3$ boolean variables), and this remains an interesting open problem.

6.2 Predicates over larger Alphabets

Our results deal only with predicates that get two input values in $\{0, 1\}$. A natural generalization is to sparsify a VCSP that uses a predicate over an alphabet of size k , i.e., $P : [k] \times [k] \rightarrow \{0, 1\}$, where $[k] := \{0, 1, \dots, k-1\}$. One predicate that we can easily sparsify is **NE** (not-equal), which is satisfied if the two constrained variables have are assigned different values. Indeed, in the graphs language, this is called a **k-Cut**, where the value of a partition (S_0, \dots, S_{k-1}) of the vertices is the total weight of all edges with endpoints in different parts. It turns out that ϵ -Cut-sparsifier is in particular an ϵ -k-Cut-sparsifier, using the following well-known double-counting argument:

$$\begin{aligned} \text{k-Cut}_{G_\epsilon}(S_0, \dots, S_{k-1}) &= \frac{1}{2} \cdot [\text{Cut}_{G_\epsilon}(S_0, \overline{S_0}) + \dots + \text{Cut}_{G_\epsilon}(S_{k-1}, \overline{S_{k-1}})] \\ &\in (1 \pm \epsilon) \cdot \frac{1}{2} \cdot [\text{Cut}_G(S_0, \overline{S_0}) + \dots + \text{Cut}_G(S_{k-1}, \overline{S_{k-1}})] \\ &= (1 \pm \epsilon) \cdot \text{k-Cut}_G(S_0, \dots, S_{k-1}) . \end{aligned}$$

In contrast, linear-equation predicates are non-sparsifiable for alphabet $[k]$ of size $k \geq 3$. Specifically, for $a \in [k]$, let the predicate Sum_a be satisfied by $x, y \in [k]$ iff $x + y = a \pmod{k}$. Then for every positively weighted digraph $G = (V, E, w)$, and every $\epsilon \in (0, 1)$, $a \in [k]$, every Sum_a - ϵ -sparsifier $G_\epsilon = (V, E_\epsilon, w_\epsilon)$ of G must have $E = E_\epsilon$. The argument is similar to the proof of [Theorem 3.2](#). Assume for contradiction there exist $e \in E \setminus E_\epsilon$. Choose $x, y, z \in [k]$ that satisfy $x + y = a$, however the three sums $z + x$, $z + y$, $z + z$ are all not equal to $a \pmod{k}$; this is clearly possible for $k \geq 4$, and easily verified by case analysis for $k = 3$. Consider an assignment where the endpoints of e have values x and y , respectively, and all other vertices have value z . Under this assignment, the value of G is $w(e) > 0$, while the value of G_ϵ is zero, a contradiction.

6.3 Sketching

In [Theorem 4.1](#) we showed that for every predicate $P \in \{\text{Cut}, \text{unCut}, \text{Or}, \text{nAnd}, \overline{10}\}$, the sketch-size of every ϵ - P -sketching-scheme is $\Omega(n/\epsilon^2)$.

Let us now address predicates with a single 1 in their truth table. In the spirit of the proof of [Theorem 3.2](#), given encoding sk_G by an ϵ -**And**-sketching-scheme we can completely restore the graph G . As there are $2^{\binom{n}{2}}$ different graphs, the sketch-size of every ϵ -**And**-sketching-scheme is at least $\Omega(n^2)$ bits. Imitating the proof of [Theorem 3.6](#), we can extend this lower bound to Dicut, 01 and 10.

For-each sketches. In order to reduce storage space of a sketch, one might weaken the requirements even further and allow the sketch to give a good approximation only with high probability. A *for-each sketching scheme* is a pair of algorithms (sk, est) ; algorithm sk is a randomized algorithm that given a graph G returns a string sk_G , whose distribution we denote by \mathcal{D}_G ; algorithm est is given such a string sk_G and a subset $S \subseteq V$, and returns (deterministically) a value $\text{est}(\text{sk}_G, S)$. We say that it is an (ϵ, δ) - P -sketching-scheme if

$$\forall G = (V, E, w), \forall S \subseteq V, \quad \Pr_{\text{sk}_G \in \mathcal{D}_G} [\text{est}(\text{sk}_G, S) \in (1 \pm \epsilon) \cdot P_G(S)] \geq 1 - \delta .$$

[\[ACK⁺15\]](#) showed that if we consider n -vertex graphs with weights only in the range $[1, W]$, then there is an $(\epsilon, 1/\text{poly}(n))$ -Cut-sketching-scheme with sketch-size $\tilde{O}(n\epsilon^{-1} \cdot \log \log W)$ bits. Imitating [Theorem 3.5](#), we can construct $(\epsilon, 1/\text{poly}(n))$ - P -sketching-scheme with the same sketch-size for every

predicate P whose truth table does not have a single 1 (and weights restricted to the range $[1, W]$). A nearly-matching lower bound by [ACK⁺15] shows that for every $\epsilon \in (2/n, 1/2)$, every $(\epsilon, 1/10)$ -Cut-sketching-scheme must have sketch-size $\Omega(n/\epsilon)$. Using γ , this lower bound can be extended to **unCut**. This technique does not work for predicates with three 1's in their truth table. Fortunately, we can duplicate the proof of [ACK⁺15] while replacing **Cut** by **Or** and using the fact that for every two vertices v, u in the graph G , it holds that $\text{Or}(\{v\}) + \text{Or}(\{u\}) - \text{Or}(\{v, u\}) = \mathbf{1}_{\{\{u, v\} \in E\}}$. We omit the details of this straightforward argument. A reduction from **OR** using γ and equation 5 will extend the lower bound also to **nAnd**, $\overline{10}$ and $\overline{01}$.

Given a sketch sk_G (i.e., one sample from distribution \mathcal{D}_G) which encodes an (ϵ, δ) -**And**-sketching-scheme, one can reconstruct every edge of G (every bit of the adjacency matrix) with constant probability. Standard information-theoretical arguments (indexing problem) imply that the sketch-size of every (ϵ, δ) -**And**-sketching-scheme is $\Omega(n^2)$ bits. Using γ we can extend this lower bound to **Dicut**, **01** and **10**.

6.4 unCut Spectral Sparsifiers

Given an undirected n -vertex graph $G = (V, E, w)$, the Laplacian matrix is defined as $L_G = D_G - A_G$ where A_G is the adjacency matrix (i.e. $A_{i,j} = w_{i,j} = w(\{v_i, v_j\})$) and D_G is a diagonal matrix of degrees (i.e. $D_{i,i} = \sum_{j \neq i} w_{i,j}$ and for $i \neq j$, $D_{i,j} = 0$). For every $x \in \mathbb{R}^n$ it holds that $x^t L_G x = \sum_{\{v_i, v_j\} \in E} w_{i,j} \cdot (x_i - x_j)^2$. In particular, for $\mathbf{1}_S$ the indicator vector of some subset $S \subseteq V$ it holds that $\mathbf{1}_S^t L_G \mathbf{1}_S = \text{Cut}_G(S)$. A subgraph H of G is called an ϵ -spectral-sparsifier of G if

$$\forall x \in \mathbb{R}^n, \quad x^t L_H x \in (1 \pm \epsilon) \cdot x^t L_G x .$$

Note that an ϵ -spectral-sparsifier is in particular an ϵ -**Cut**-sparsifier. Nonetheless, spectral sparsifiers preserve additional properties such as the eigenvalues of the Laplacian matrix (approximately). [BSS14] showed that every graph admits an ϵ -spectral-sparsifier with $O(n/\epsilon^2)$ edges.

Definition 6.3. *Given a graph G , we call $U_G = (D_G + A_G)$ the Negated Laplacian of G . Given a subset $S \subseteq V$, let $\phi_S \in \mathbb{R}^n$ be a vector such that $\phi_{S,i} = 1$ if $v_i \in S$ and $\phi_{S,i} = -1$ otherwise.*

One can verify that for arbitrary $x \in \mathbb{R}^n$,

$$x^t U_G x = \sum_{i < j} w_{i,j} \cdot (x_i + x_j)^2$$

In particular, for every subset $S \subseteq V$, it holds that

$$\phi_S^t U_G \phi_S = 4 \cdot \text{unCut}_G(S) .$$

Next, we will show how we can use U_G to construct an **unCut**-sparsifier G_ϵ (in alternative way to Theorem 3.5) such that U_{G_ϵ} has (approximately) the same eigenvalues as U_G . A matrix $M \in \mathbb{R}^{n \times n}$ is called *BSDD* (*Balanced Symmetric Diagonally Dominant*) if $M = M^t$ and for every index i , $M_{i,i} = \sum_{j \neq i} |M_{i,j}|$. Note that L_G and U_G are both BSDD. A matrix M' is *governed* by M if whenever $M'_{i,j} \neq 0$, also $M_{i,j} \neq 0$ and has the same sign. Note that if H is a subgraph of G then U_H is governed by U_G . A matrix M' is called an ϵ -spectral-sparsifier of M if M' is governed by M and

$$\forall x \in \mathbb{R}^n, \quad x^t M' x \in (1 \pm \epsilon) \cdot x^t M x .$$

The following was implicitly shown in [ACK⁺15].

Theorem 6.4 ([ACK⁺15]). *Given BSDD matrix $M \in \mathbb{R}^{n \times n}$ and parameter $\epsilon \in (0, 1)$, there is an ϵ -spectral-sparsifier M' for M where M' is BSDD matrix with $O(n/\epsilon^2)$ non-zero entries.*

Fix a graph G and parameter ϵ , according to Theorem 6.4, there is a BSDD balanced matrix H with $O(n/\epsilon^2)$ non-zero entries, that governed by U_G which is a ϵ -spectral-sparsifier for U_G . All this properties define a unique graph G_ϵ such that $U_{G_\epsilon} = H$. In particular G_ϵ is ϵ -unCut-sparsifier of G with $O(n/\epsilon^2)$ edges.

References

- [ACK⁺15] A. Andoni, J. Chen, R. Krauthgamer, B. Qin, D. P. Woodruff, and Q. Zhang. On sketching quadratic forms. Preprint, earlier versions are available as [arXiv:1403.7058](#) and [arXiv:1412.8225](#), April 2015.
- [AG09] K. J. Ahn and S. Guha. Graph sparsification in the semi-streaming model. In *36th International Colloquium on Automata, Languages and Programming, ICALP '09*, pages 328–338. Springer-Verlag, 2009. [arXiv:0902.0140](#), [doi:10.1007/978-3-642-02930-1_27](#).
- [BHM80] R. A. Brualdi, F. Harary, and Z. Miller. Bigraphs versus digraphs via matrices. *J. Graph Theory*, 4(1):51–73, 1980. [doi:10.1002/jgt.3190040107](#).
- [BK96] A. A. Benczúr and D. R. Karger. Approximating s-t minimum cuts in $\tilde{O}(n^2)$ time. In *28th Annual ACM Symposium on Theory of Computing*, pages 47–55. ACM, 1996. [doi:10.1145/237814.237827](#).
- [BK02] A. A. Benczúr and D. R. Karger. Randomized approximation schemes for cuts and flows in capacitated graphs. *CoRR*, cs.DS/0207078, 2002. [arXiv:cs/0207078](#).
- [BSS14] J. D. Batson, D. A. Spielman, and N. Srivastava. Twice-ramanujan sparsifiers. *SIAM Review*, 56(2):315–334, 2014. [doi:10.1137/130949117](#).
- [dCHS11] M. K. de Carli Silva, N. J. A. Harvey, and C. M. Sato. Sparse sums of positive semidefinite matrices. *CoRR*, abs/1107.0088, 2011. [arXiv:1107.0088](#).
- [FHHP11] W. S. Fung, R. Hariharan, N. J. Harvey, and D. Panigrahi. A general framework for graph sparsification. In *43rd Annual ACM Symposium on Theory of Computing*, pages 71–80. ACM, 2011. [doi:10.1145/1993636.1993647](#).
- [KK15] D. Kogan and R. Krauthgamer. Sketching cuts in graphs and hypergraphs. In *Conference on Innovations in Theoretical Computer Science*, pages 367–376. ACM, 2015. [doi:10.1145/2688073.2688093](#).
- [KL02] D. R. Karger and M. S. Levine. Random sampling in residual graphs. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 63–66, 2002.
- [KP12] M. Kapralov and R. Panigrahy. Spectral sparsification via random spanners. In *3rd Innovations in Theoretical Computer Science Conference*, pages 393–398. ACM, 2012. [doi:10.1145/2090236.2090267](#).

- [Mad10] A. Madry. Fast approximation algorithms for cut-based problems in undirected graphs. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 245–254. IEEE, 2010.
- [NR13] I. Newman and Y. Rabinovich. On multiplicative λ -approximations and some geometric applications. *SIAM Journal on Computing*, 42(3):855–883, 2013. doi:[10.1137/100801809](https://doi.org/10.1137/100801809).
- [She09] J. Sherman. Breaking the multicommodity flow barrier for $O(\sqrt{\log n})$ -approximations to sparsest cut. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 363–372, 2009.
- [SS11] D. A. Spielman and N. Srivastava. Graph sparsification by effective resistances. *SIAM J. Comput.*, 40(6):1913–1926, December 2011. doi:[10.1137/080734029](https://doi.org/10.1137/080734029).
- [ST04] D. A. Spielman and S.-H. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *36th Annual ACM Symposium on Theory of Computing*, pages 81–90. ACM, 2004. doi:[10.1145/1007352.1007372](https://doi.org/10.1145/1007352.1007372).
- [ST11] D. A. Spielman and S.-H. Teng. Spectral sparsification of graphs. *SIAM J. Comput.*, 40(4):981–1025, July 2011. doi:[10.1137/08074489X](https://doi.org/10.1137/08074489X).